



Special Report: Internet Filtering Alternatives

**A White Paper
For IT Managers, CTOs and HR Directors**

Table of Contents

Executive Summary	1
Internet Use at Work: Legal and Productivity Issues	2
Statistics on Internet Abuse.....	3
Selecting an Internet Filtering Solution	4
Internet Access Control Filtering Methodologies.....	4
Keyword Filtering.....	4
URL-based Filtering Versus IP Address Filtering.....	5
Importance of Maintaining the Database	5
Going Beyond Web-based Content	6
Internet Access Management Filtering Solutions.....	6
Software-based Approaches.....	6
Client-based Filtering Software	6
Server-based Web Filtering Software	7
Add-ons to Firewalls.....	8
Hardware-based Approaches	9
Turnkey Filtering Servers	9
Conclusion.....	11
Appendix A: Comparison of Internet Filtering Solutions.....	i

Executive Summary

In the Internet-dependent world of today, it is difficult to imagine a workplace without access to the Web. Yet, much of the information available to employees on the Internet is not job related. What started as a productivity boon has gradually turned into a bandwidth and productivity drain with huge potential legal liability.

In response, organizations have increasingly sought ways to proactively control Web access. This demand has driven the development of a variety of Web-filtering methodologies. However, the growth of new Internet content, combined with the need for simple network installation and straightforward ways to effectively manage large user communities, has made most of these alternatives too cumbersome.

This paper describes the most prevalent Internet filtering technologies, with a focus on a new generation of appliance-based products that specifically address management, reliability and performance issues. These appliances are scalable, easy to install and administer, and offer customizable user profiles and filtering mechanisms of all Internet content types. As a result, the appliance approach empowers organizations to take full control of their Internet access issues.

Internet Use at Work: Legal and Productivity Issues

In a 2002 study by the Computer Security Institute (CSI), 78 percent of polled enterprises reported employee abuse of Internet access privileges, to include downloading pirated software or pornography, shopping on the Internet and inappropriate use of email systems.

Indiscriminate Web surfing by employees carries serious legal risks. Bringing sexually oriented material into the workplace (whether from the Internet or from other sources) has the potential to create a hostile work environment, which, under federal and state sex discrimination laws, can be grounds for a sexual harassment lawsuit. Downloading pirated software and music also has legal implications as do Web content copyright infringement activities and Internet gambling.

Education and government face strict regulations that require Web filtering. If filtering is not in place, government funding may be forfeited. The Children's Internet Protection Act (CIPA) requires schools and libraries to block visual depiction of pornography, obscenity or other offensive material to children in order to qualify for government funds that help pay for computers and Internet services.

While legal liability is often the impetus for an organization to initiate Internet monitoring, it is not the only risk associated with Web surfing. According to recent statistics, inappropriate Web surfing also puts employee productivity at tremendous risk. A Gallup poll found that the average employee spends 75 minutes per workday surfing the Net. Another recent survey found that 72 percent of the employees polled admitted to reading the news online on a regular basis, 45 percent said they made travel arrangements and 40 percent made regular online purchases. At an hourly employee rate of \$20 per hour (the Department of Labor's estimate of the average cost of employing an American worker), companies are losing an average of \$125 per week per worker. Even more is lost when professional employees spend time surfing the Net inappropriately.

In addition, there are technological drains and costs related to Internet use. In the same survey, 13 percent of the employees interviewed took advantage of their company's high-speed Internet connections to download music while on the job. Thanks to popular applications such as streaming audio and video, Internet-based games and downloadable music files, frivolous use of the Internet can create severe bandwidth problems. This abuse can be even worse at colleges and universities. One employee's inappropriate use may negatively affect another employee's speed of access or storage space available for work-related files. Such abuse can be quite expensive, given the cost of the network access and hardware necessary to accommodate increased network traffic and data storage.

Statistics on Internet Abuse

Statistics on Internet Abuse

- Sex site surfing is reported in 62% of organizations. (*PC Week*)
- Internet surfing on the job accounts for 30-40% of lost worker productivity. (IDC Research)
- 70% of all Internet porn traffic occurs during the 9-to-5 workday. (SexTracker)
- 32.6% of workers have no specific objective when they surf the Internet. (eMarketer.com)
- One in five men and one in eight women admitted using their work computers as their primary lifeline to access sexually explicit material online. (MSNBC)
- Web users at the office take advantage of high-speed connections to access broadband entertainment sites such as Broadcast.com and MP3.com more frequently than at home. (Nielsen/Net Ratings)
- 82% of U.S. business executives surveyed by the consulting firm Dataquest (a division of the Gartner Group) believe Internet use should be monitored at their companies. (InformationWeek Online)

Selecting an Internet Filtering Solution

There are many types of Web filtering tools from which to choose. The following factors enter into the decision; each is discussed in the following sections of this document.

- The type of filtering methodology used
- How often the database used for blocking Internet requests is updated and how it is maintained
- Other types of Internet services that can be included in the blocking process
- Whether to employ a software-based filtering product or a dedicated hardware appliance

Internet Access Control Filtering Methodologies

The primary goal of Internet access management is to identify and block or monitor objectionable or inappropriate content. Several filtering methods -- including filtering by keyword, URL or IP address -- have been developed toward this goal. Since each method identifies objectionable or inappropriate content in a slightly different manner, some are more effective than others in meeting the delicate balancing act of filtering out too little or too much material.

Keyword Filtering

The simplest method of Web filtering is to block by keyword. Keyword filtering scans for specific words within the text of a page as it is downloaded. The page is blocked if any one of the listed words is detected.

The major downfall of keyword blocking is that it cannot take context into account and, as a result, can too often block acceptable content. For instance, keyword filtering on the word "breast" may inadvertently block sites that contain valuable research information on "breast cancer." Because keyword filtering works by scanning only the text portion of the requested Web sites, it also is of little value for blocking pages that contain nothing but non-textual photographs (such as those found on pornography sites). In addition, keyword blocking can be slow, which impacts user response time while accessing the Internet.

URL-based Filtering Versus IP Address Filtering

More sophisticated filtering methods employ a database of URL or Internet Protocol (IP) address information to block access to specific, predetermined sites. Blocking all traffic from a specific IP address has the advantage of being simple and fast. However, in today's environment where a number of different sites might be "virtually hosted" at a shared IP address, the IP filtering method blocks either *all* of the virtually hosted sites or *none*. In this scenario, blocking one inappropriate site on a hosted server would preclude access to a dozen other hosted sites with acceptable content.

URL blocking, on the other hand, can provide blocking down to specific pages within a Web site. Using this approach, a generic photo archive site containing mostly acceptable content, with only a few pages dedicated to nudity, could be made accessible with only the objectionable pages blocked.

Today's more sophisticated Internet access management techniques, such as those found in St. Bernard's iPrism appliance, generally combine extensive URL filtering along with the ability to handle direct entry of IP addresses.

Importance of Maintaining the Database

Thousands of new Web sites are added to the Internet every day. As a result, one of the core criteria in selecting an Internet filtering solution is how accurate and current the filtering database is. No matter what method of blocking is used, any filtering scheme must include mechanisms for continuous updates and maintenance of the database used to make blocking decisions. Even the most comprehensive database will rapidly become outdated without constant review and updating.

In addition, structuring the database into logical categories with mechanisms for adding customer-defined categories increases the value of the filtering product, because these functions enable organizations to tailor their filtering policies to meet their specific requirements.

The importance of human review in the database maintenance process cannot be understated. Using automated tools to identify new Web sites that *may* contain unacceptable content is a great place to start. However, relying solely on those tools for making the final determination can lead to many inaccuracies that end up annoying some users desiring access to appropriate sites and access to objectionable sites for others.

Going Beyond Web-based Content

Another consideration in selecting a filtering solution is whether it addresses the significant amount of non-Web content that is also available on the Internet. As the variety and types of Internet content have proliferated, the issue of access control has expanded well beyond just HTML-based Web content. Access to Internet services such as ICQ® and IRC chat, FTP downloads, RealAudio® broadcasts and MP3 music can easily consume significant bandwidth and resources while degrading employee productivity. Web filtering products that address these types of content and services add even more value to access control.

Internet Access Management Filtering Solutions

There are many types of Internet access management tools to choose from. In addition to deciding whether to choose a solution that filters Internet usage based on keywords, URLs, or IP addresses, organizations must consider whether to take a software or hardware approach.

Software-based Approaches

Client-based Filtering Software

The first wave of Web filtering products primarily used a client-side model in which the filtering engine itself resided on the desktop computer.

These desktop Web filtering products operate in a similar fashion to anti-virus programs that monitor the local machine's activities. These products block Internet content based upon the current listings in a local database, which must be updated periodically at each individual PC.

Client-side deployments have become quite popular for consumer-oriented Web filtering, because they are cost effective for very small numbers of users. However, this approach has a number of shortcomings when used in larger organizations:

- **High administration costs.** The coordination and effort involved with updating many different desktop-resident databases can present a difficult challenge in large environments, resulting in outdated and/or inconsistent filtering.
- **Potentially unreliable.** Since Web filtering software resides on each user's computer, it is more vulnerable to tampering or circumvention. In addition, many of

these products use keyword filtering that may block appropriate sites or not block inappropriate sites, as previously discussed.

- **Lack of flexibility.** Because the filtering software is tied to each individual machine, it is difficult to provide special permissions for roving individuals (such as teachers) or different permissions for multiple people who share a single PC.
- **Performance degradation.** Local installation of software filtering software may also slow down the local PC due to drains on memory, storage and processing resources.
- **Compatibility issues.** Lastly, not all PCs on a heterogeneous network may be compatible with the client software.

Server-based Web Filtering Software

Server-based Web filtering software addresses many of the limitations of client-based filtering software. Instead of installing the filtering software on each local PC, the filtering engine is typically installed on an existing application server (NT or Unix) configured either as a proxy server or as part of a firewall.

The server-based filtering engine automatically intercepts all packets requesting external Web-based content. It screens and filters all inappropriate content for the entire network from a unified central location using a single database.

This approach is less prone to tampering by individual users. It is also much simpler to administer, because there is only one database to update. Existing server hardware *may* be able to be used, so this solution may have lower costs upfront, as compared to other solutions, such as turnkey server solutions and appliances (discussed later in this section).

However, a server-based filtering solution has hidden costs and limitations, including:

- **High administration and support costs.** Server-based filtering solutions can require a significant amount of IT administration time to install, test, configure and maintain, especially as new versions or patches of the add-on software or server operating system are released. Even after installation has been successfully completed and the server tuned for acceptable performance, any changes or upgrades to the computing environment can require re-configuration or re-tuning of the Web filtering server. The operating system and proxy must be properly configured for optimal performance. If there are problems, server-based filtering

software is more difficult to support, with multiple vendors (server hardware, server OS, firewall, etc.) required for problem resolution.

- **Potential security risks.** Since multiple vendors may need to become involved for hands-on support, network security may be at risk. In contrast, the troubleshooting of dedicated filtering appliances does not require system administration level access to an enterprise's server, which reduces the risk of security breaches.
- **Performance degradation.** Server-based filtering software uses server resources, which can degrade network performance. Organizations may need to purchase a separate server just for filtering. Depending on the superiority of the server, this solution usually offers lower performance, less scalability and higher costs than dedicated filtering appliances.
- **Compatibility issues.** Depending upon the myriad of variables in a typical network-computing environment (such as Windows NT, Solaris, Novell, proxy servers, firewalls, etc.) the filtering server may require different configuration parameters and in some cases may not even be compatible with existing systems. Upgrades and new versions of operating systems or firewalls may require patches or new versions of the filtering software, leading to an endless cycle of upgrades -- or in some cases, eventual abandonment of the software due to incompatibility.

Add-ons to Firewalls

Filtering services can also be added to firewalls or other network devices, allowing companies to use existing hardware if performance, capacity and security considerations permit. For this reason, add-ons may be less expensive initially than dedicated filtering appliances or turnkey filtering server solutions.

However, firewalls are typically not optimized for large Web site database lookups, so performance and scalability are often issues, in addition to other limitations. There have also been situations reported in the media where filtering software added to firewalls has compromised the security provided by the firewall.

Most of the pros and cons associated with the add-on server software filtering approach apply to this approach as well.

Hardware-based Approaches

Turnkey Filtering Servers

Turnkey filtering solutions combine general-purpose server hardware with pre-configured filtering software, resulting in higher upfront costs as compared to server software-based solutions. However, they are easier to install and support and are typically more scalable and higher performing when compared to both client- and server-software based filtering.

However, when compared to dedicated filtering appliances, turnkey server solutions have several limitations:

- **Setup and support are more difficult than a dedicated appliance.** Turnkey filtering servers require more initial setup effort than dedicated filtering appliances. They are also more difficult to support and may require extensive knowledge of NT, Linux or Unix in order to install and troubleshoot.
- **Reliability, scalability and performance are not as high as a dedicated appliance.** Because turnkey filtering solutions use a general-purpose operating system added to general-purpose hardware, they will be less reliable and robust than an appliance that has been optimized for filtering reliability and performance.
- **Larger footprint than most dedicated appliances.** Turnkey solutions use general-purpose server hardware, which often takes up more space than a dedicated filtering appliance. Combined with the other limitations, turnkey server solutions may not be appropriate for very small organizations.

Dedicated Filtering Appliances

Similar to the shift that has occurred in the firewall and router markets from using software to using hardware-based solutions, organizations are moving toward implementing dedicated appliances to meet their Internet access management needs.

According to the Frost & Sullivan 2001 *Content Filtering Markets* report, “Network administrators dread the implementation of a content filtering solution for their networks because of the work required to install, maintain, update and monitor the solution ... For this reason, filtering appliances are beginning to proliferate, as they are a solution that requires minimal setup time. The future of the corporate content market will increasingly demand filtering appliances ...”

Dedicated filtering appliances are easier to implement than other solutions because they can be installed into an existing network with little effort and impact. They are optimized for Internet filtering, so they also tend to offer the highest performance and scalability. Another benefit is that there is a single point of contact for service and support. Overall, dedicated filtering appliances offer the best solution with the lowest total cost of ownership (TCO) when compared to the other solutions.

Appliances offer the following specific advantages:

- **Lower costs and administration overall.** No software to install, test or manage, and no additional software or hardware to purchase.
- **Automatic database upgrades and software updates** are normally included with a dedicated appliance, as compared to software vendors who only provide automatic database updates.
- There is only **one vendor to call for support**; plus hardware and software maintenance is normally included.
- **Lower TCO** through simpler initial and ongoing administration. Appliances are generally easier to install and maintain, resulting in IT administration time that can be used for other purposes.
- Appliances are platform independent and can **adapt to any environment**. Companies can change network architecture or add or upgrade resources such as servers, routers, firewalls, or workstation operating systems without having to change appliances.
- **True remote management**, which is not usually the case with most software approaches to Internet filtering.

Disadvantages of this approach include:

- **Higher upfront costs when compared to some software add-on solutions.** For this reason, a dedicated filtering appliance may not be the most cost effective solution for very small organizations, such as those with less than 25 Internet connected users.
- **Demo software is not available for testing.** However, some companies, such as St. Bernard, may offer evaluation units prior to purchase or provide online demo capabilities.

Conclusion

Performance, reliability, ease of use and administration, and total cost of ownership are factors to consider when selecting an Internet filtering solution. In general, appliances are superior to software-based approaches to Internet filtering because they get the job done reliably, quickly and with the lowest total cost of ownership.

Other factors besides the software versus hardware issue include:

- The type of methodology used to filter Web pages. Keyword filtering and IP address filtering are the simplest methods for vendors to implement, however URL-based filtering provides greater accuracy so users aren't blocked from appropriate sites.
- The comprehensiveness and accuracy of the database used to filter. Organizations will want to select Internet filtering vendors that use human reviewers and take great care to ensure that their databases do not filter too much or too little. Databases need to be updated frequently and customizable at the customer level.
- The types of other Web services that can also be managed by the filtering solution, for example IRC chat sessions and MP3 downloads.

The Appendix summarizes the advantages and disadvantages of each approach described in this paper.

An additional white paper on St. Bernard's dedicated Internet filtering appliance, iPrism, can be found at <http://www.stbernard.com/iprism>.

Appendix A: Comparison of Internet Filtering Solutions

Hardware-based Solutions		
Type of Filtering Solution	Pluses	Minuses
<p>Dedicated Filtering Appliances</p> <p>This category includes hardware devices specifically designed for Internet access monitoring and control.</p> <p><i>Example: iPrism by St. Bernard Software</i></p>	<p>Easiest alternative to implement, since it's a self-contained appliance that can be installed into an existing network without impacting existing servers, firewalls or other network resources.</p> <p>A dedicated device designed specifically for Internet monitoring and filtering, offering the highest performance and scalability of all solutions.</p> <p>Lower total cost of ownership (TCO) than with network software or client software-based solutions.</p> <p>Rack Mountable.</p> <p>Single point of contact for service and support.</p>	<p>Upfront costs may be higher than software-only solutions even though the total cost of ownership (TCO) is lower.</p> <p>May not be the most cost effective solution for very small organizations.</p> <p>Cannot download demo software to test. (However, with iPrism, 30-day evaluation units are available upon request, as is a live demo.)</p>
<p>Turnkey Filtering Servers</p> <p>This category includes solutions where general-purpose server hardware is bundled with pre-configured filtering software.</p> <p><i>Example: 8e6 R2000</i></p>	<p>Easier installation than software add-on only or client-based solutions, although not as easy as dedicated filtering appliances.</p> <p>Better TCO than with software add-on or client-based solutions, but may not be as favorable as dedicated filtering appliances.</p> <p>May be a more scalable and higher-performance solution than software add-on solutions, but typically not as much as dedicated filtering appliances.</p> <p>Single point of contact for service and support.</p>	<p>Upfront costs may be higher than software add-on solutions.</p> <p>May not be appropriate for very small organizations.</p> <p>Requires more initial setup effort, and often expense, than dedicated filtering appliances.</p> <p>Uses general-purpose operating system (OS) added to general-purpose hardware, which will typically be less reliable than dedicated filtering appliances.</p> <p>More difficult to support than filtering server appliances.</p> <p>Larger physical size than most dedicated filtering appliances, and may not be rack mountable.</p> <p>Some solutions may require modification of each desktop browser's proxy setting to provide filtered Internet access.</p>

Software-based Solutions

Type of Filtering Solution	Pluses	Minuses
<p>Server-based Filtering Software</p> <p>This category includes software that works on NT or Unix operating system platforms, or as a plug-in for proxy server software offerings.</p> <p><i>Examples: Websense Proxy Server (standalone on NT) and Websense for Microsoft Proxy Server by Websense, I-Gear for NT or Solaris by Symantec, SurfControl for Microsoft Proxy Server by SurfControl, Internet Manager by Elron Software</i></p>	<p>May be able to use existing server hardware if available, and performance, capacity and security considerations permit.</p> <p>There may be less upfront costs than with dedicated filtering appliance or turnkey filtering server solutions.</p> <p>Initial testing usually easier, software can often be downloaded from the Internet.</p>	<p>More implementation effort required than with dedicated filtering appliance or turnkey filtering server solutions: initial software installation, compatibility testing, and associated production system downtime.</p> <p>Additional installation, compatibility testing, etc. required as patches or new versions of the add-on software are released. Also, patches and new versions of the firewall software may be incompatible with the add-on filtering software and require it to be upgraded or patched as well (and those upgrades may not be available in the same time frame that the firewall software is).</p> <p>Utilizes server resources: CPU cycles, memory and hard disk. Organizations often need to purchase a separate server just for filtering.</p> <p>Typically lower performance and less scalable than dedicated filtering appliances.</p> <p>Potential reliability issues, OS and proxy must be properly configured and optimized for best performance.</p> <p>More difficult to support. Since the filtering vendor does not typically provide hardware support, multiple vendors may be required for problem resolution.</p> <p>Vendor must have system administration level access in order to access the server to provide direct hands-on support, thus potentially compromising network security. With dedicated filtering appliances, vendors are only accessing their own system.</p> <p>Configuration dependency based on OS or proxy type and version may exist.</p> <p>Upgrades to either OS or proxy may impact compatibility and usability of filtering software.</p> <p>Some solutions may require modification of individual desktop browsers to provide filtered Internet access.</p>

Type of Filtering Solution	Pluses	Minuses
<p>Add-ons to Firewalls</p> <p>Filtering services that can be added on to firewalls or other network devices.</p> <p><i>Examples: Websense for CheckPoint FireWall-1 by Websense, SurfControl for Firewall -1</i></p>	<p>Can use existing network hardware if available, and performance, capacity and security considerations permit.</p> <p>There may be less upfront costs than with a dedicated filtering server appliance and turnkey filtering server solutions.</p>	<p>Firewalls typically are not optimized for large Web site database lookups, so performance and scalability can be an issue, particularly when compared to dedicated filtering appliances.</p> <p>Disk and memory resources typically limited, not scalable to keep up with rapid growth in number of filtered Internet sites.</p> <p>Using third party application add-on may impact security of firewall (example: CyberPatrol creating a security breach in Network Associates' Gauntlet firewall).</p> <p>Must be very careful with configuration changes so as not to impact firewall security, operation and performance.</p> <p>More implementation effort required than with dedicated filtering appliance or turnkey filtering server solutions: initial software installation, compatibility testing, and associated production system downtime.</p>
<p>Client-based Filtering Software</p> <p>This category includes Windows software, Mac software and browser plug-ins that add Web filtering capabilities to a desktop PC.</p> <p><i>Examples: CyberPatrol for Windows and Mac by SurfControl, Net Nanny by Net Nanny Software Inc.</i></p>	<p>For organizations with very small numbers of users, it can be the most cost effective option.</p> <p>Perceived as simple solution.</p> <p>Easy to test - just download, install on a desktop PC and go.</p>	<p>If necessary to load client software on a lot of PCs in company or school environments, is very time consuming for administrators and intrusive to users.</p> <p>Can be relatively easy for workers and kids to circumvent protection.</p> <p>Many use keyword filtering techniques that are unreliable, either not blocking inappropriate sites or blocking access to useful and appropriate sites.</p> <p>Uses memory, disk space and processor resources on every client PC.</p> <p>Filtering database updates must applied, usually manually, to each PC.</p> <p>Client software may not be compatible with all PCs in the network.</p> <p>Primary targets are home users, not business or schools.</p>